arXiv:1609.03045v1 [stat.ME] 10 Sep 2016

PRINCIPAL COMPONENT ANALYSIS AND THE LOCUS OF THE FRÉCHET MEAN IN THE SPACE OF PHYLOGENETIC TREES

TOM M. W. NYE, XIAOXIAN TANG, GRADY WEYENBERG, AND RURIKO YOSHIDA

Abstract. Most biological data are multidimensional, posing a major challenge to human comprehension and computational analysis. Principal component analysis is the most popular approach to rendering two- or three-dimensional representations of the major trends in such multidimensional data. The problem of multidimensionality is acute in the rapidly growing area of phylogenomics. Evolutionary relationships are represented by phylogenetic trees, and very typically a phylogenomic analysis results in a collection of such trees, one for each gene in the analysis. Principal component analysis offers a means of quantifying variation and summarizing a collection of phylogenies by dimensional reduction. However, the space of all possible phylogenies on a fixed set of species does not form a Euclidean vector space, so principal component analysis must be reformulated in the geometry of tree-space, which is a CAT(0) geodesic metric space. Previous work has focused on construction of the first principal component, or principal geodesic. Here we propose a geometric object which represents a k-th order principal component: the locus of the weighted Fréchet mean of $k+1$ points in tree-space, where the weights vary over the standard k-dimensional simplex. We establish basic properties of these objects, in particular that locally they generically have dimension k, and we propose an efficient algorithm for projection onto these surfaces. Combined with a stochastic optimization algorithm, this projection algorithm gives a procedure for constructing a principal component of arbitrary order in tree-space. Simulation studies confirm these algorithms perform well, and they are applied to data sets of Apicomplexa gene trees and the African coelacanth genome. The results enable visualizations of slices of tree-space, revealing structure within these complex data sets.

1. Introduction

One of the great opportunities offered by modern genomics is that phylogenetics applied on a genomic scale (phylogenomics) should be especially powerful for elucidating gene and genome evolution, relationships among species and populations, and processes of speciation and molecular evolution. However, a well-recognized hurdle is the sheer volume of genomic data that can now be generated relatively cheaply and quickly, but for which analytical tools are lagging. There is a major need to explore new approaches to undertake comparative genomic and phylogenomic studies much more rapidly and robustly than existing tools allow. Principal component analysis is the most popular approach for reducing the dimension of multidimensional data sets. The problem of multidimensionality is acute in the rapidly growing area of phylogenomics, which can provide insight into relationships and evolutionary patterns of a diversity of organisms, from humans, plants and animals, to microbes and viruses.

Data sets consisting of collections of phylogenetic trees are challenging to analyse, due both to high dimensionality and the complexity of the space containing the data. Multivariate statistical procedures such as outlier detection [37], clustering [12] and multidimensional scaling [16] have previously been applied to such data sets. However, principal component analysis is perhaps the most useful multivariate statistical tool for exploring high-dimensional data sets, due to its flexibility and its attractive properties. For example, [39] and [7] recently showed that principal component analysis automatically projects to the subspace where the global solution of K-means clustering lies, and so facilitates K-means clustering to find near-optimal solutions. Although principal component analysis for data in $\mathbb{R}^m$ can be defined in several different ways, the following description is natural for reformulating the procedure in tree-space. Suppose we have data $Z = \{z_1, z_2, \ldots, z_n\}$ where $z_i \in \mathbb{R}^m$ for $i = 1, \ldots, n$. For any set of $k+1$ points $V = \{v_0, v_1, \ldots, v_k\} \subset \mathbb{R}^m$ we can define

$$\Pi(V) = \left\{ \sum_{i=0}^{k} p_i v_i : p_0, \ldots, p_k \in \mathbb{R}, p_0 + \cdots + p_k = 1 \right\} \tag{1}$$

so that $\Pi(V)$ is the hyperplane in $\mathbb{R}^m$ containing $v_0, v_1, \ldots, v_k$. The orthogonal L^2 distance of any point $y \in \mathbb{R}^m$ from $\Pi(V)$ is denoted $d(y, \Pi(V))$, and the sum of squared projected distances of the data Z onto the hyperplane is defined by

$$D_Z^2\left(\Pi(V)\right) = \sum_{i=1}^{n} d(z_i, \Pi(V))^2.$$

Then the k-th order principal component Π_k corresponds to a choice of V which *minimizes* this sum. In $\mathbb{R}^m$, Π_0 is the sample mean, Π_1 is the line through the sample mean which minimizes the sum of squared projected distances, and so on for $k = 2, 3, \ldots$. Although it is not explicit in the definition above, in $\mathbb{R}^m$ the principal components are nested:

$$\Pi_0 \subset \Pi_1 \subset \Pi_2 \subset \cdots . \tag{2}$$

This description of principal component analysis relies heavily on the vector space properties of $\mathbb{R}^m$: $\Pi(V)$ is defined as a linear combination of vectors and the procedure uses orthogonal projection.

However, the space of phylogenetic trees with $N+1$ leaves is not an Euclidean vector space. It follows that we cannot directly apply classical principal component analysis to a data set consisting of phylogenetic trees. The set $\mathcal{T}_N$ of all phylogenetic trees with $N+1$ leaves labelled $0, 1, \ldots, N$ is a so-called CAT(0) space [4, 6]. This means that $\mathcal{T}_N$ is a metric space with a unique geodesic (shortest length path) between any pair of points, with the geodesic computable in $O(N^4)$ steps [29]. Amongst other properties, projection onto closed sets is well-defined in CAT(0) spaces. The analogue of the zero-th order principal component is given by the *Fréchet mean* of the data $z_1, \ldots, z_n$. The Fréchet mean is a statistic which characterizes the central tendency of a distribution in arbitrary metric spaces. For any metric space S equipped with metric $d(\cdot, \cdot)$, the Fréchet population mean, μ, with respect to distribution ν is defined by

$$\mu(\nu) = \arg\min_{y \in S} \int_S d(y, x)^2 d\nu(x).$$

The discrete analogue, the weighted Fréchet mean of a sample $Z = \{z_1, \ldots, z_n\}$ with respect to a weight vector w, is

$$\mu(Z, w) = \arg\min_{y \in S} \sum_{i=1}^{n} w_i \, d(y, z_i)^2,$$

where the weights w_i satisfy $w_i \geq 0$ for $i = 1, \ldots, n$. In any CAT(0) space, $\mu(Z, w)$ is a well-defined unique point given data Z and weight vector w. The definition of the zero-th order principal component Π_0 in $\mathbb{R}^m$ given above coincides with the definition of the Fréchet sample mean with weights $w_i = 1$ in any CAT(0) space. Several algorithms for computing the Fréchet sample mean in $\mathcal{T}_N$ have been developed [3, 26] and we review these later in Section 2.2 as they play an important role in our methodology.

Methods for constructing a principal geodesic in tree-space, an analogue of $\Pi_1 \subset \mathbb{R}^m$ as defined above, have recently been developed. In [28], the approach involved firing geodesics from some mean tree. For each candidate geodesic Γ, the sum of squared projected distances $D_Z^2(\Gamma)$ was computed and a greedy algorithm was used to adjust Γ in order to mimimize $D_Z^2(\Gamma)$. The geodesics considered were infinitely long, but these have the disadvantage that in some cases many such geodesics fit the data equally well. Subsequent approaches therefore considered finitely long geodesic segments [9, 27]. The geodesic segment between two points $v_0, v_1 \in \mathcal{T}_N$ is analogous to $\Pi(V)$ in equation (1) with $k = 1$, except that the weights p_0, p_1 must constrained to be positive under the analogy. [9] constrained the ends of the geodesic to be points in the sample Z and sought the corresponding geodesic Γ which mimimizes $D_Z^2(\Gamma)$, whereas [27] did not restrict the geodesic and used a stochastic optimization algorithm to perform the minimization.

In this paper we address two fundamental questions: (i) which geometric object most naturally plays the role of a k-th order principal component in tree-space; and (ii) given such an object, how can we efficiently project data points onto the object? Our proposed solution is to replace the definition of $\Pi(V) \subset \mathbb{R}^m$ given in equation (1) with the locus of the weighted Fréchet mean of points $v_0, v_1, \ldots, v_k$ in tree-space. Specifically, suppose $V = \{v_0, v_1, \ldots, v_k : v_i \in \mathcal{T}_N, i = 0, 1, \ldots, k\}$ and define $\Pi(V) \subset \mathcal{T}_N$ by

$$\Pi(V) = \{\mu(V, p) : p \in \mathcal{S}^k\}$$

where $\mathcal{S}^k$ is the k-dimensional simplex of probability vectors

$$\mathcal{S}^k = \{(p_0, p_1, \ldots, p_k) : p_i \geq 0, i = 0, 1, \ldots, k \text{ and } \sum_i p_i = 1\}$$

and $\mu(V, p)$ is the Fréchet mean of the points in set V with weights p. We call $\Pi(V)$ the *locus of the Fréchet mean* of V. Our choice of notation is intended to emphasize the analogy between the definition of $\Pi(V)$ in tree-space and the corresponding definition for $\mathbb{R}^m$ in equation (1). The locus of the Fréchet mean is a type of *minimal surface*, as the following physical analogy suggests: imagine connecting a point $y \in \mathcal{T}_N$ to points $v_0, v_1, \ldots, v_k \in \mathcal{T}_N$ by $k + 1$ pieces of elastic. When the point y is free to move, it will move under the action of the elastic into an equilibrium position in tree-space. We can imagine how this equilibrium point changes as the stiffness in the pieces of elastic is varied, which corresponds to varying $p \in \mathcal{S}^k$. As the equilibrium point moves around it scans out a surface in tree-space. In Euclidean space the locus of the Fréchet mean of some collection

of points is a subset of a hyperplane. In tree-space, as we will show, the hypersurface can be curved. Surfaces of this kind have recently been studied by [31] in the context of Riemannian manifolds and other geodesic metric spaces. We discuss the relationship of the present paper to that work in Section 6.

Our main theoretical results are as follows. First, when $V = \{v_0, v_1, \ldots, v_k\}$ we derive a set of local implicit equations for $\Pi(V)$. These allow us to derive conditions for $\Pi(V)$ to be locally flat, and also enable us to construct explicit realizations of $\Pi(V)$ in certain interesting cases. Secondly, using the implicit equations we show the locus of the Fréchet mean $\Pi(V)$ in $\mathcal{T}_N$ is locally k dimensional for generic V, and thus forms a suitable candidate for a k-th order principal component. Third, we present an algorithm for projection onto $\Pi(V)$ which relies only on the CAT(0) properties of $\mathcal{T}_N$. We demonstrate accuracy of the projection algorithm via a simulation study.

The remainder of the paper is organized in the following way. Section 2 reviews the fundamental concepts and results in phylogenetic tree-space: the construction of geodesics, calculation of the Fréchet mean and convex hulls. Section 3 studies the geometric features of the locus of the Fréchet mean, in particular establishing its dimension with Theorem 1. We draw the reader's attention to Section 3.4 which contains an explicit example which illustrates geodesics in tree-space, convex hulls, and details of the calculation of dimension. Section 4 presents the algorithm for projecting sets of phylogenetic tree data onto the locus of the Fréchet mean of $k+1$ fixed trees and describes the algorithm used to fit these objects to the data. This section also contains a simulation study to test the effectiveness of the algorithms. In Section 5 we apply these methods to two data sets. The first is a set of gene trees coming from fish and tetrapods intended to investigate the relationship of coelacanth and lungfish to the tetrapods, and the second is a set of gene trees from apicomplexa. We make some concluding remarks in Section 6.

2. The geometry of tree-space

2.1. **Construction of tree-space and its geodesics.** Throughout the paper, the m-dimensional Euclidean vector space is denoted by $\mathbb{R}^m$. The non-negative and positive orthants in $\mathbb{R}^m$ are denoted by $\mathbb{R}^m_{\geq 0}$ and $\mathbb{R}^m_{>0}$, respectively. For any vectors $x, y \in \mathbb{R}^m$, $||x||$ denotes the Euclidean norm of x, and $\langle x, y\rangle$ denotes the Euclidean inner product.

A *phylogenetic tree* with the leaf set $X = \{0, 1, \ldots, N\}$ is an undirected weighted acyclic graph with $N+1$ degree-1 vertices labelled $0, 1, \ldots, N$, and with no degree-2 verices. We consider rooted trees, and the root is the leaf labelled "0". Each such tree contains $N + 1$ *pendant* edges, which connect to the leaves, and up to $N - 2$ *internal edges*. The maximum number of internal edges is achieved when the tree is binary, in which case all vertices have degree 3 other than the leaves, in which case the tree is called *fully resolved*. If a tree contains fewer edges then it is called *unresolved* and there must be at least one vertex with degree ≥ 4. Each edge in a phylogeny is assigned a strictly positive weight (also called the edge *length*). Given a tree $x \in \mathcal{T}_N$, the set of edges of x is denoted $\mathcal{E}(x)$, and the weight associated to $e \in \mathcal{E}(x)$ is denoted $|e|_x$. It is convenient to define $|e|_x$ to be zero whenever e is not contained in x. Tree-space $\mathcal{T}_N$ is the set of all phylogenetic trees with leaf set X [4].

Tree-space can be embedded in $\mathbb{R}^M$ for $M = 2^N - 1$ in the following way. If we cut any edge $e \in \mathcal{E}(x)$ then the tree x splits into two disconnected pieces. This determines a *split* $X_e|X_e^c$ of the leaf set X, where $X_e \cup X_e^c = X$ and $X_e \cap X_e^c = \emptyset$. By convention we choose X_e to be the set containing the root 0, and so there are $M = 2^N - 1$ possible splits of X. The collection of splits represented by a tree x is called the *topology* of x. Since edges and splits are equivalent, we use the notation $\mathcal{E}(x)$ to also represent the set of splits in x. By choosing some arbitrary ordering of the set of all splits, each tree $x \in \mathcal{T}_N$ can be represented as a vector in $\mathbb{R}^M$ with up to $2N-1$ positive entries given by the edge weights of x, and zeros for each split which is not contained in x. However, an arbitrary choice of vector will not necessarily represent a tree: for example the splits $\{0,1\}|\{2,3,\ldots,N\}$ and $\{0,2\}|\{1,3,\ldots,N\}$ cannot both be contained in the same tree so any vector for which these splits both have strictly positive value does not represent a tree. Two splits $X_e|X_e^c, X_f|X_f^c$ are *compatible* if one of the four sets $X_e \cap X_f$, $X_e^c \cap X_f$, $X_e \cap X_f^c$, $X_e^c \cap X_f^c$ is empty, in which case there is at least one tree containing both splits. Any collection of pairwise compatible splits determines a valid tree topology [33, Theorem 3.1.4].

The embedding into Euclidean space reveals the combinatorial structure of $\mathcal{T}_N$. Every tree $x \in \mathcal{T}_N$ contains N pendant edges and so $\mathcal{T}_N$ is the product of $\mathbb{R}^N_{>0}$ and a space corresponding to the internal edges. It is therefore convenient to ignore the pendant edges, and consider the corresponding embedding of tree-space into $\mathcal{R}_N = \mathbb{R}^{M-N}$. Given any tree topology τ containing m internal edges, the set of trees with topology τ corresponds to a subset $\mathcal{O}_\tau \subset \mathcal{R}_N$ which is isomorphic to $\mathbb{R}^m_{>0}$. (The isomorphism is with respect to the local Euclidean structure.) Each such region is called the *orthant* for topology τ. The boundary of $\mathcal{O}_\tau$ in $\mathcal{R}_N$ corresponds to trees obtained by removing one or more internal edges from τ. Equivalently, the trees on the boundary can be obtained by taking a tree x in $\mathcal{O}_\tau$ and continuously shrinking one or more internal edges down to length zero. Thus, for a fully-resolved topology τ, the codimension-1 boundaries of $\mathcal{O}_\tau$ correspond to trees containing $N-3$ internal edges, and in general each codimension-k boundary corresponds to trees containing $N-k-2$ internal edges, for $k = 1,\ldots,N-2$. There are $(2N-3)!!$ possible fully resolved rooted tree topologies and so $\mathcal{T}_N$ is built from $(2N-3)!!$ orthants isomorphic to $\mathbb{R}^{N-2}_{>0}$ together with the boundaries of these orthants which correspond to trees which are not fully resolved. Orthants are glued together at their boundaries since a given unresolved tree containing m internal edges can be obtained by removing edges from several different trees containing $m+1$ edges. Orthants corresponding to fully-resolved topologies are glued at their codimension-1 boundaries in a relatively simple way. If a single internal edge in a tree with fully-resolved topology τ is contracted to length zero and removed from the tree, the result is a vertex of degree 4. There are 3 possible ways to add in an additional edge to give another fully resolved topology (including the original edge which was removed) so each codimension-1 face of $\mathcal{O}_\tau$ is glued to two other such orthants. Trees containing no internal edges are called *star-trees*: the point $0 \in \mathcal{R}_N$ corresponds to the set of star-trees and is contained in the boundary of every orthant $\mathcal{O}_\tau$.

The topology of $\mathcal{T}_N$ is taken to be that induced by the embedding into Euclidean space. Geodesics are constructed by considering continous paths in $\mathcal{T}_N$ which are Euclidean straight-line segments in each orthant. The length of a path is the sum of the Euclidean segment lengths. As shown in [4], the shortest such path or *geodesic*

between two points $x, y \in \mathcal{T}_N$ is unique, and it will be denoted $\Gamma(x, y)$. The distance $d(x, y)$ is defined to be the length of $\Gamma(x, y)$ and this defines the metric $d(\cdot, \cdot)$ on $\mathcal{T}_N$. By definition, $d(x, y)$ incorporates information about both the topologies and edge lengths of x and y. Given two points x, y in the same orthant $\Gamma(x, y)$ is simply the Euclidean line segment between x, y, whereas when x, y are in different orthants $\Gamma(x, y)$ consists of a series of straight line segments traversing orthants corresponding to different topologies. [4] proved that $\mathcal{T}_N$ is a CAT(0) space, and so it has several additional geometrical properties [6].

[29] established a polynomial time ($O(N^4)$) algorithm to compute the geodesic between any two trees in $\mathcal{T}_N$. The details of their algorithm are not important for the present application, but we do require some notation for the form of the geodesics it constructs. Given $x, y \in \mathcal{T}_N$ let $\mathcal{C}(x, y)$ be the set of splits in $\mathcal{E}(x) \cup \mathcal{E}(y)$ which are compatible with every split in $\mathcal{E}(x)$ and every split in $\mathcal{E}(y)$. Adopting notation from [29], the geodesic $\Gamma(x, y)$ is characterized by disjoint sets of internal splits

$$A_{xy}^{(1)}, A_{xy}^{(2)}, \ldots, A_{xy}^{(\ell_{xy})} \subset \mathcal{E}(x), \quad \text{and}$$
$$B_{xy}^{(1)}, B_{xy}^{(2)}, \ldots, B_{xy}^{(\ell_{xy})} \subset \mathcal{E}(y)$$

where $\ell_{xy} \geq 0$ is an integer which depends on x, y. These sets of splits determine the order in which edges are removed and added as the geodesic is traversed. The union $\bigcup A_{xy}^{(j)}$ for $j = 1 \ldots, \ell_{xy}$ is $\mathcal{E}(x) \setminus \mathcal{C}(x, y)$ and similarly for tree y. We let $\mathcal{A}(x, y)$ be the ordered list of sets $(A_{xy}^{(j)} : j = 1, \ldots, \ell_{xy})$ and similarly define $\mathcal{B}(x, y)$. The *support* of $\Gamma(x, y)$ is defined to be the triple $(\mathcal{A}(x, y), \mathcal{B}(x, y), \mathcal{C}(x, y))$. It characterizes the sequence of orthants the geodesic traverses. For any set $E \subset \mathcal{E}(x)$ we adopt the notation

$$\|E\|_x = \left(\sum_{e \in E} |e|_x^2 \right)^{1/2}$$

and similarly for subsets of $\mathcal{E}(y)$. [29] showed that

$$d(x, y)^2 = \|A_{xy} + B_{xy}\|^2 + \|C_{xy} - D_{xy}\|^2 \tag{3}$$

where A_{xy} is the ℓ_{xy}-dimensional vector whose j-th element is $\|A_{xy}^{(j)}\|_x$ and similarly for B_{xy} the j-th element is $\|B_{xy}^{(j)}\|_y$. Vectors C_{xy} and D_{xy} have dimension $|\mathcal{C}(x, y)|$, and respectively contain the edge lengths $|e|_x$ and $|e|_y$ for $e \in \mathcal{C}(x, y)$. It follows from equation (3) that

$$d(x, y)^2 = \|x\|^2 + \|y\|^2 + 2\langle A_{xy}, B_{xy} \rangle - 2\langle C_{xy}, D_{xy} \rangle \tag{4}$$

where $\|x\|^2$ is the sum of squared edge lengths in x and similarly for y.

The following definition characterizes certain geodesics which behave rather like Euclidean straight lines.

Definition 1. *Suppose $x, y \in \mathcal{T}_N$ are fully resolved. The geodesic $\Gamma(x, y)$ is called* simple *if each set $A_{xy}^{(i)}$ and $B_{xy}^{(i)}$ contains exactly one element for $i = 1, \ldots, \ell_{xy}$. Equivalently, $\Gamma(x, y)$ is simple if and only if at most one edge length contracts to zero at a time as the geodesic is traversed.*

The following definition determines the set of trees y such that the geodesics $\Gamma(x, y)$ to a fixed point x all share the same support.

Definition 2. *Fix some point $x \in \mathcal{T}_N$ and an orthant $\mathcal{O}_\tau$ corresponding to a fully-resolved topology τ. Given any valid support σ, the set*

$$S_x(\sigma, \tau) = \{y \in \mathcal{O}_\tau : \Gamma(x, y) \text{ has support } \sigma\}$$

is called a support region.

[26] considered very similar subsets of $\mathcal{T}_N$ and established their properties. Given x and τ there are only finitely many possible supports. We will use the fact that the union over the set of possible supports

$$\bigcup_\sigma S_x^\circ(\sigma, \tau)$$

is dense in $\mathcal{O}_\tau$, where $S_x^\circ(\sigma, \tau)$ denotes the interior of each support region. The boundaries between the support regions are continuous codimension-1 surfaces within each orthant.

2.2. **Algorithms for computing the Frechét mean.** Several algorithms for computing the unweighted or weighted Fréchet mean of a sample in $\mathcal{T}_N$ have been developed [34, 26, 3]. These algorithms have the following general structure. Let the sample of trees be $Z = \{z_1, \ldots, z_n\} \subset \mathcal{T}_N$. At the i-th iteration there is an estimate μ_i of the Fréchet mean of Z. To find the next estimate, μ_{i+1}, a data point z_j is selected, either deterministically or stochastically depending on the particular algorithm. The geodesic $\Gamma(\mu_i, z_j)$ is constructed, and μ_{i+1} is taken to be the point a certain proportion of the distance along the geodesic. This proportion can depend on the weights when the weighted Fréchet mean is estimated. In each case, some form of convergence of the sequence $\mu_0, \mu_1, \mu_2, \ldots$ to the Fréchet mean of Z can be proved, independent of the initial estimate μ_0.

Our methodology does not make direct use of these algorithms. However, our proposed algorithm for projecting data onto the locus of the Fréchet mean is adapted from the algorithm of [34] (see Section 4.1), and so we present Sturm's algorithm here. The algorithm computes the Fréchet mean of $z_1, \ldots, z_n$ using weights $p_1, \ldots, p_n \geq 0$. By definition, the Fréchet mean is invariant under positive scaling of the weights, so we can take $\sum p_i = 1$. Sturm's algorithm proceeds in the following way.

Algorithm 1. *Sturm's algorithm for the weighted Fréchet mean.*
Fix an initial estimate μ_0 and set $i = 0$.
***Repeat**:*

1. *Sample $Z_i \in \{z_1, \ldots, z_n\}$ such that $\Pr(Z_i = z_j) = p_j$.*
2. *Construct $\Gamma(\mu_i, Z_i)$.*
3. *Let μ_{i+1} be the point a proportion s_i along $\Gamma(\mu_i, Z_i)$ where $s_i = 1/(i+2)$.*
4. *Set $i \leftarrow i + 1$.*

***Until** the sequence $\mu_0, \mu_1, \ldots$ converges.*

Convergence can be tested in various ways, for example repeating until a specified number of consecutive estimates μ_i all lie within distance ϵ of each other. Sturm proved that the points μ_i converge in probability to the Fréchet mean of the distribution defined by sampling $z_1, \ldots, z_n$ according to probabilities $p_1, \ldots, p_n$.

The deterministic algorithm of [3] for computing the weighted Fréchet mean is similar to Sturm's algorithm, except the data points are used cyclically, as opposed to being randomly sampled as the algorithm progresses, and the weighting is instead

taken into account in the definition of the proportions s_i. We use the algorithm of [3] for computing the Fréchet mean in order to test our projection algorithm, and this procedure is also described in Section 4.1.

2.3. **Convex hulls.** [27] suggested that the convex hull of $k+1$ points in $\mathcal{T}_N$ might be a suitable geometrical object to represent a k-th order principal component. A set $A \subset \mathcal{T}_N$ is *convex* if and only if for all points $x, y \in A$ the geodesic $\Gamma(x, y)$ is also contained in A. The *convex hull* of a set of points is the smallest convex set containing those points. Any geodesic segment is the convex hull of its end-points, and using the convex hull of 3 points to represent a second order principal component is a natural generalization of the idea of a principal geodesic. Convexity is also a desirable property when performing projections, as occurs in a principal component analysis. However, convex hulls in tree-space do not have the correct dimension. Examples for which the convex hull of 3 points is 3 dimensional can readily be constructed [38, 24]. It is demonstrated in [23] that the dimension of a convex hull of 3 points in $\mathcal{T}_N$ can be arbitrarily high as N increases. More generally, convex hulls in tree-space are difficult to characterize geometrically with several fundamental questions unanswered. These issues make convex hulls less appealing as geometrical objects to represent principal components, and at this point we turn our attention to the locus of the Fréchet mean. However we demonstrate the relationship between the locus of the Fréchet mean and convex hull for an explicit configuration of 3 points $v_0, v_1, v_2 \in \mathcal{T}_N$ later in Section 3.4.

3. The locus of the Fréchet mean

3.1. **Basic properties.** Throughout this section we work with $k+1$ points $v_0, v_1, \ldots, v_k \in \mathcal{T}_N$ and let $V = \{v_0, v_1, \ldots, v_k\}$. As in the Introduction, we define $\mu : (\mathcal{T}_N)^{k+1} \times \mathcal{S}^k \to \mathcal{T}_N$ by

$$\mu(V, p) = \arg\min_{x \in \mathcal{T}_N} \sum_{i=0}^{k} p_i \, d(x, v_i)^2.$$

The locus of the Fréchet mean of V, denoted $\Pi(V) \subset \mathcal{T}_N$, is

$$\Pi(V) = \{\mu(V, p) : p \in \mathcal{S}^k\}.$$

Here we establish some basic properties of $\Pi(V)$, while the next section presents a more detailed analysis of $\Pi(V)$ within orthant interiors. First, the map μ is continuous and so $\Pi(V)$ is compact since it is the continuous image of a compact set. Continuity of μ can be proved using the deterministic algorithm for calculating the weighted Fréchet mean given by [3]: the output of the algorithm depends continuously on the inputs V and p. (We do not give a detailed proof for reasons of brevity.) Secondly, the points $v_0, \ldots, v_k$ are contained in $\Pi(V)$, since $\mu(V, e_i) = v_i$ where e_i denotes the i-th standard basis vector in $\mathcal{S}^k$. Similarly each geodesic $\Gamma(v_i, v_j)$ is contained in $\Pi(V)$, by taking p to be a convex combination of e_i and e_j. By the same argument, $\Pi(V)$ contains $\Pi(W)$ where W is any non-empty subset of V.

In Euclidean space the convex hull of $k+1$ points coincides with the locus of the Fréchet mean of the points. However, this is not the case in tree-space, though $\Pi(V)$ is contained in the closure of the convex hull of V. This follows because any point in $\Pi(V)$ can be approximated arbitrarily closely by performing a finite number of steps of the algorithm of [3] (see Section 2.2). Provided the algorithm is initialized

with one of the points $v_0, \ldots, v_k$, each of these steps remains within the convex hull, and so the limit point is contained in the closure of the convex hull. It is important to note that $\Pi(V)$ is itself generally *not* convex. As a consequence, there might not be a unique closest point on $\Pi(V)$ to any given point z, although the minimum distance of z from $\Pi(V)$ is well-defined. By using $\Pi(V)$ as a principal component we have therefore lost the desirable property of uniqueness of projection.

Fréchet means in tree-space exhibit a property called *stickiness* [17]. This essentially means that for fixed V the map $\mu(V, \cdot) : \mathcal{S}^k \to \mathcal{T}_N$ can fail to be injective. Specifically, depending on the points in V, there might exist open sets in $\mathcal{S}^k$ which all map to the same point in tree-space. This has implications when we project data points onto $\Pi(V)$: given a data point z, the value of p which minimizes $d(z, \mu(V, p))^2$ might be non-unique, even if there is a unique closest point $x \in \Pi(V)$ to z.

3.2. **Implicit equations for the locus of the Fréchet mean.** The algebraic form of tree-space geodesics described in Section 2.1 can be used to derive implicit equations for the edge lengths of trees lying on the locus of the Fréchet mean $\Pi(V)$, and these equations are fundamental to establishing the dimension of $\Pi(V)$. For fixed $V = \{v_0, \ldots, v_k\}$ consider the objective function $\Omega : \mathcal{T}_N \times \mathcal{S}^k \to \mathbb{R}$ defined by

$$\Omega(x, p) = \sum_{i=0}^{k} p_i \, d(x, v_i)^2.$$

Suppose we fix an orthant $\mathcal{O}_\tau$ for a fully resolved topology τ. Let $x \in \mathcal{O}_\tau$ have edge lengths $x_j = |e_j|_x$ where $e_j \in \mathcal{E}(x)$ for $j = 1, \ldots, 2N-1$. [26] showed that functions of the form $d(x, v_i)^2$ are continuously differentiable on $\mathcal{O}_\tau$ with respect to the edge lengths x_j. In order to minimize Ω we additionally assume x lies in a set

$$S = \bigcap_{i=0}^{k} S^\circ_{v_i}(\sigma_i, \tau)$$

for some choice of supports $\sigma_0, \ldots, \sigma_k$. We call sets of this form *mutual support regions* with respect to $v_0, \ldots, v_k$. They are dense in $\mathcal{O}_\tau$ using the properties of support regions given in Section 2.1. Each mutual support region is essentially a piece of tree-space for which the combinatorics of the geodesics to $v_0, \ldots, v_k$ do not vary as a reference point moves around the region. An example of a decomposition of orthants into mutual support regions is given in Section 3.4. Under this assumption on x, we can write down the algebraic form of $d(x, v_i)^2$ using equation (4) to give

$$\Omega(x, p) = \|x\|^2 + \sum_{i=0}^{k} p_i \, \left(\|v_i\|^2 + 2\langle A_{xv_i}, B_{xv_i} \rangle - 2\langle C_{xv_i}, D_{xv_i} \rangle \right)$$

so

$$\frac{\partial \Omega}{\partial x_j} = 2x_j + 2 \sum_{i=0}^{k} p_i \, \frac{\partial}{\partial x_j} \left(\langle A_{xv_i}, B_{xv_i} \rangle - \langle C_{xv_i}, D_{xv_i} \rangle \right). \tag{5}$$

If the point $x \in S$ lies on the locus of the Fréchet mean $\Pi(V)$ then $\partial\Omega/\partial x_j = 0$ for all j, and so we want to evaluate these derivatives to obtain implicit equations relating the edge lengths x_j to the vector p.

Let y be any of the trees $v_0, \ldots, v_k$. By definition

$$\langle C_{xy}, D_{xy} \rangle = \sum_{e \in \mathcal{C}(x,y)} |e|_x |e|_y$$

so

$$\frac{\partial}{\partial x_j} \langle C_{xy}, D_{xy} \rangle = |e_j|_y$$

since x_j is the length of split e_j and so the derivative of $\langle C_{xy}, D_{xy} \rangle$ is just a constant. The term $\langle A_{xy}, B_{xy} \rangle$ has a more general functional dependence on x_j. By definition

$$\begin{aligned} \langle A_{xy}, B_{xy} \rangle &= \sum_{l=1}^{\ell_{xy}} \|A_{xy}^{(l)}\|_x \|B_{xy}^{(l)}\|_y \\ &= \sum_{l=1}^{\ell_{xy}} \left(\sum_{e \in A_{xv_i}^{(l)}} |e|_x^2 \right)^{1/2} \left(\sum_{f \in B_{xy}^{(l)}} |f|_y^2 \right)^{1/2}. \end{aligned}$$

For any edge $e_j \in \mathcal{C}(x,y)$ this expression does not depend on x_j so the derivative is zero. When $e_j \in \mathcal{E}(x) \setminus \mathcal{C}(x,y)$ only the first term in brackets will depend on x_j. Since the sets $A_{xy}^{(l)}$ are disjoint it must be the case that e_j is contained in exactly one set and we define r_{ij} to be the index l of that set when $y = v_i$. Then

$$\frac{\partial}{\partial x_j} \langle A_{xv_i}, B_{xv_i} \rangle = \|B_{xv_i}^{(r_{ij})}\| \frac{\partial}{\partial x_j} \left(\sum_{e \in A_{xv_i}^{(r_{ij})}} |e|_x^2 \right)^{1/2} = x_j \frac{\|B_{xv_i}^{(r_{ij})}\|}{\|A_{xv_i}^{(r_{ij})}\|}.$$

In the case that $A_{xv_i}^{(r_{ij})}$ contains only e_j and no other splits, we have $\|A_{xv_i}^{(r_{ij})}\| = x_j$ so the expression becomes

$$\frac{\partial}{\partial x_j} \langle A_{xv_i}, B_{xv_i} \rangle = \|B_{xv_i}^{(r_{ij})}\|$$

which is a constant. Substituting these expressions into equation (5) gives

$$\frac{\partial \Omega}{\partial x_j} = 2x_j + 2 \sum_{i=0}^{k} p_i \left(x_j \frac{\|B_{xv_i}^{(r_{ij})}\|}{\|A_{xv_i}^{(r_{ij})}\|} (1 - \mathcal{C}_{ij}) - |e_j|_{v_i} \mathcal{C}_{ij} \right) \tag{6}$$

where $\mathcal{C}_{ij} = 1$ if $e_j \in \mathcal{C}(x, v_i)$ and is zero otherwise.

We define $F : \mathcal{O}_\tau \times \mathcal{S}^k \to \mathbb{R}^{2N-1}$ by

$$F(x,p) = \nabla_x \Omega(x,p). \tag{7}$$

This function is continuously differentiable with respect to the edge lengths for all x lying within the interior of mutual support regions. On the boundary between mutual support regions F is continuous but may not be differentiable. In section 3.3 we show the matrix of second derivatives of Ω is positive definite on each mutual support region, and so every solution to $\nabla_x \Omega = 0$ is a minimum. It follows that $\Pi(V)$ is locally the solution to $F(x,p) = 0$.

The following lemma establishes conditions for $\Pi(V)$ to be a hyperplane within the mutual support region $S \subset \mathcal{O}_\tau$.

Lemma 1. *If the supports $\sigma_0, \ldots, \sigma_k$ are such that the geodesics $\Gamma(x, v_i)$ are simple for all $i = 0, \ldots, k$ (in the sense of definition 1) then $\Pi(V)$ is a hyperplane of dimension k or lower in $S = \bigcap_i S^\circ_{v_i}(\sigma_i, \tau)$.*

Proof. If all the geodesics $\Gamma(x, v_i)$ are simple for $x \in S$ then each set $A^{(l)}_{xv_i}$ contains exactly one split. Then equation (6) becomes

$$\frac{\partial \Omega}{\partial x_j} = 2x_j + 2\sum_{i=0}^{k} p_i \alpha_{ij}$$

for some constants α_{ij}. Solving $F(x, p) = 0$ gives each edge length x_j as a linear combination of $p_0, p_1, \ldots, p_k$, which establishes the result. Generically, $\Pi(V)$ is therefore locally a k-dimensional hyperplane, but the dimension may be lower. Further discussion about this point is given in Section 3.3. □

3.3. **The dimension of the locus of the Fréchet mean.** We aim to prove that $\Pi(V)$ has dimension k in each mutual support region. The strategy is to first show that the matrix of second derivatives of Ω, or equivalently the matrix with elements $\partial F_j / \partial x_k$ with F defined by equation (7), is positive definite. Calculation of the dimension of $\Pi(V)$ follows by applying the implicit function theorem.

Lemma 2. *The matrix with elements $\partial F_j / \partial x_k$ is positive definite for all x in the mutual support region S.*

Proof. Using equation (6) we have

$$\frac{\partial F_j}{\partial x_k} = 2\delta_{jk} + 2\sum_{i=0}^{k} p_i Q^{(i)}_{jk}$$

where the matrix $Q^{(i)}$ has elements

$$Q^{(i)}_{jk} = \frac{\partial}{\partial_k}\left(x_j \frac{\|B^{(r_{ij})}_{xv_i}\|}{\|A^{(r_{ij})}_{xv_i}\|}(1 - \mathcal{C}_{ij}) - |e_j|_{v_i}\mathcal{C}_{ij} \right).$$

We start by assuming $\mathcal{C}_{ij} = 0$ for all i, j (so that $e_i \notin \mathcal{C}(x, v_i)$ for all i, j), and drop this assumption later. Then

$$\begin{aligned} Q^{(i)}_{jk} &= \delta_{jk}\frac{\|B^{(r_{ij})}_{xv_i}\|}{\|A^{(r_{ij})}_{xv_i}\|} - x_j \frac{\|B^{(r_{ij})}_{xv_i}\|}{\|A^{(r_{ij})}_{xv_i}\|^2}\frac{\partial}{\partial x_j}\|A^{(r_{ij})}_{xv_i}\| \\ &= \delta_{jk}\frac{\|B^{(r_{ij})}_{xv_i}\|}{\|A^{(r_{ij})}_{xv_i}\|} - x_j x_k I^{(i)}_{jk}\frac{\|B^{(r_{ij})}_{xv_i}\|}{\|A^{(r_{ij})}_{xv_i}\|^3} \end{aligned}$$

where $I^{(i)}$ is a $(2N-1) \times (2N-1)$ dimensional matrix with $I^{(i)}_{jk} = 1$ whenever e_k is contained in $A^{(r_{ij})}_{xv_i}$ and zero otherwise. Equivalently $I^{(i)}$ indicates whether splits e_j and e_k are simultaneously contracted to zero on $\Gamma(x, v_i)$. We will show that the matrices $Q^{(i)}$ are positive semi-definite. For any vector $\xi \in \mathbb{R}^{2N-1}$ we have

$$\sum_{j,k} Q^{(i)}_{jk}\xi_j\xi_k = \sum_j \frac{b_j}{a_j}\xi_j^2 - \sum_{j,k}\xi_j\xi_k x_j x_k \frac{b_j}{a_j^3} I^{(i)}_{jk} \tag{8}$$

where $a_j = \|A^{(r_{ij})}_{xv_i}\|$ and $b_j = \|B^{(r_{ij})}_{xv_i}\|$. Now fix a single set of splits $A^{(l)}_{xv_i}$ and let J_l denote the indices of splits in this set. If we restrict the right-hand side of the

last equation to indices $j \in J_l$ we obtain

$$\sum_{j\in J_l} \frac{b_j}{a_j}\xi_j^2 - \sum_{j,k\in J_l} \xi_j\xi_k x_j x_k \frac{b_j}{a_j^3}.$$

The terms a_j adopt the same value for all $j \in J_l$, and similarly for b_j, so they are independent of the summation index in the last expression. Also for $j \in J_l$

$$a_j^2 = \|A^{(l)}_{xv_i}\|_x = \sum_{m\in J_l} x_m^2.$$

Then

$$\sum_{j,k\in J_l} Q^{(i)}_{j,k}\xi_j\xi_k = \sum_{j\in J_l} \frac{b_j}{a_j}\xi_j^2 - \frac{\sum_{j,k\in J_l}\xi_j\xi_k x_j x_k}{\sum_{m\in J_l} x_m^2}\frac{b_j}{a_j}.$$

The Cauchy-Schwartz inequality shows the right-hand side is ≥ 0 since a_j, b_j are constant over this range of j. It follows that the right-hand side of equation (8) is ≥ 0, so each matrix $Q^{(i)}$ is positive semi-definite. If we drop the assumption that $\mathcal{C}_{ij} = 0$ for all i, j, this introduces rows and columns of zeros into each matrix $Q^{(i)}$. However, the matrices $Q^{(i)}$ must therefore remain positive semi-definite, and this establishes the calim in the statement of the proof. □

Theorem 1. *Within the mutual support region S, the locus of the Fréchet mean $\Pi(V)$ is a submanifold of dimension k or lower. For generic selections of the points $v_0, \ldots, v_k$ the dimension is k.*

Proof. Application of the implicit function theorem to the map F when $x \in S$ establishes that there is a locally-defined function $g : \mathcal{S}^k \to S$ such that $F(g(p), p) = 0$ and that the locus $(g(p), p)$ is a k-dimensional submanifold of $S \times \mathcal{S}^k$. In fact, the image $g(p) \subset S$ will be k-dimensional when the derivative of F with respect to p, $\nabla_p F$, has rank k which is the case for generic arrangements of points $v_0, \ldots, v_k$ in tree-space. This is analogous to considering the hyperplane containing $k + 1$ given points in Euclidean space: generically the hyperplane has dimension k but the dimension can be lower. □

3.4. **Explicit calculation.** In this section we construct an explicit example of the locus of the Fréchet mean for three points in $\mathcal{T}_5$. This example helps demonstrate the nature of geodesics in tree-space, the derivation of the impicit equations for $\Pi(V)$, the relationship with the convex hull and other geometrical features. We start by fixing v_0, v_1, v_2 to have the topologies and edge lengths shown in Figure 1. We will ignore the pendant edge lengths, and so the orthants containing these trees can be identified with three orthants in $\mathbb{R}^3$ equipped with standard coordinates ξ_1, ξ_2, ξ_3. There are five splits contained in these trees (other than the pendant splits): they will be denoted $\{0, 1\}$, $\{2, 3\}$, $\{4, 5\}$, $\{3, 4, 5\}$, $\{2, 3, 4\}$ by neglecting the complements in $X = \{0, 1, \ldots, N\}$. We then use the notation $x(\{0, 1\})$ to denote the length associated to split $\{0, 1\}$ in tree x, for example. Under the identification with $\mathbb{R}^3$ we have

$$\xi_1 = x(\{2,3\}) \text{ when } \{2,3\} \in x, \qquad \xi_1 = -x(\{3,4,5\}) \text{ when } \{3,4,5\} \in x,$$
$$\xi_2 = x(\{4,5\}) \text{ when } \{4,5\} \in x, \qquad \xi_2 = -x(\{2,3,4\}) \text{ when } \{2,3,4\} \in x,$$

and $\xi_3 = x(\{0, 1\})$. Figure 1 shows the location of trees v_0, v_1, v_2 under this identification. The orthant $\xi_1 < 0, \xi_2 < 0, \xi_3 > 0$ does not correspond to a valid tree topology as $\{3, 4, 5\}$ is not compatible with $\{2, 3, 4\}$. At each codimension-1 face

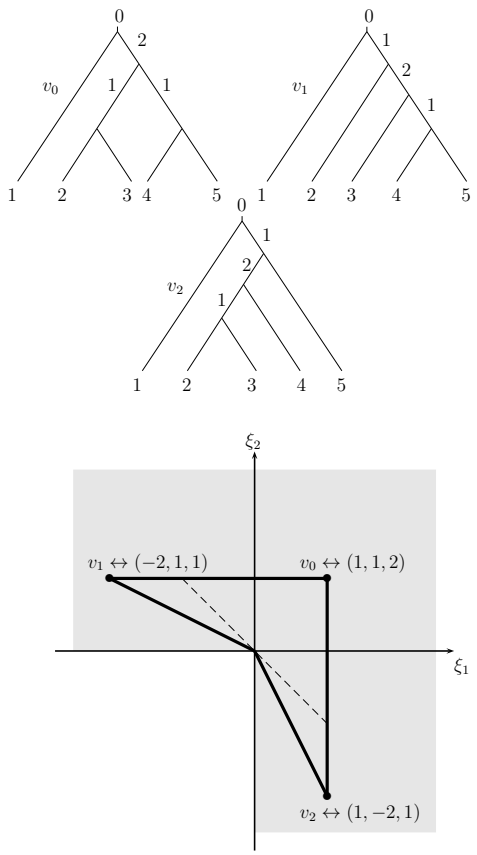

FIGURE 1. Left: topologies for the trees v_0, v_1, v_2 for the example in Section 3.4. Weights for internal edges are shown. Right: coordinates of the trees v_0, v_1, v_2 under the identification with orthants in $\mathbb{R}^3$. The ξ_3 axis points out of the page. The geodesics between v_0, v_1, v_2 are shown: $\Gamma(v_1, v_2)$ 'kinks' around the origin. The dashed line is between points $(-1, 1, 4/3)$ and $(1, -1, 4/3)$ on $\Gamma(v_0, v_1)$ and $\Gamma(v_0, v_2)$ respectively.

between the orthants shown there is in fact a third orthant in $\mathcal{T}_5$ glued at the same boundary, but these do not play a role in this example.

Figure 1 shows that the geodesics $\Gamma(v_0, v_1)$ and $\Gamma(v_0, v_2)$ are straight line segments under the identification with $\mathbb{R}^3$, while the geodesic $\Gamma(v_1, v_2)$ 'kinks' at a codimension-2 face. This behaviour is typical of geodesics in $\mathcal{T}_N$: they are straight line segents within each orthant but they can contain kinks at the voundaries between orthants. Figure 1 also shows how the convex hull of v_0, v_1, v_2 has dimension 3. The dashed line shows the geodesic between points $(-1, 1, 4/3)$ and $(1, -1, 4/3)$ on $\Gamma(v_0, v_1)$ and $\Gamma(v_0, v_2)$ respectively. The convex hull therefore contains the points $(0, 0, 1)$ and $(0, 0, 4/3)$ and so there are 4 points which are not coplanar within each orthant of the convex hull.

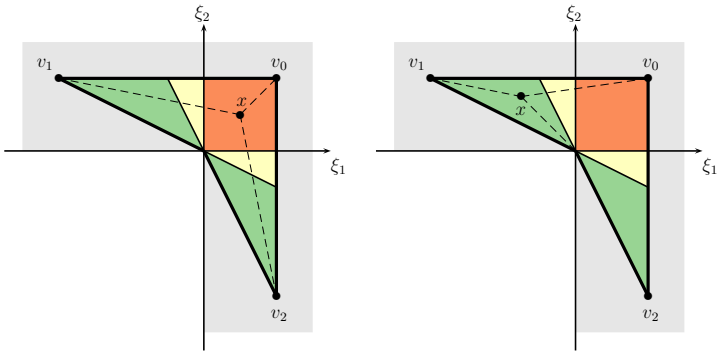

FIGURE 2. Decomposition of the locus of the Fréchet mean into mutual support regions. There are five such regions, highlighted in different colours. The dashed lines show the geodesics between a point x and the points v_0, v_1, v_2. Left: when x is contained in the orange region, none of the geodesics $\Gamma(x, v_i)$ hit codimension-2 orthant faces, and so Lemma 1 shows $\Pi(V)$ is planar. The same applies to the two yellow mutual support regions. Right: when x is contained one of the green regions then $\Gamma(x, v_2)$ is not simple (it hits a codimension-2 boundary) and so $\Pi(V)$ is not planar.

Figure 2 shows the decomposition of the orthants into mutual support regions for v_0, v_1, v_2. There are five regions in total, and the geodesics $\Gamma(x, v_i)$ are simple for all $i = 0, 1, 2$ when x is contained in three of the regions. Lemma 1 shows that $\Pi(V)$ is therefore planar in those regions with equation

$$\xi = (p_0 - 2p_1 + p_2, p_0 + p_1 - 2p_0, 1 + p_0).$$

We can also explicitly calculate equations for $\Pi(V)$ in the mutual support region contained in $2\xi_1 + \xi_2 < 0$ and shown on the right in figure 2. For x contained in this region, the squared distances to the vertices are

$$\begin{aligned} d(x, v_0)^2 &= (1 - \xi_1)^2 + (1 - \xi_2)^2 + (2 - \xi_3)^2 \\ d(x, v_1)^2 &= (2 + \xi_1)^2 + (1 - \xi_2)^2 + (1 - \xi_3)^2 \\ d(x, v_2)^2 &= \left(\sqrt{5} + (\xi_1^2 + \xi_2^2)^{1/2}\right)^2 + (1 - \xi_3)^2 \end{aligned}$$

where x has coordinates ξ_1, ξ_2, ξ_3. These can be used to write down an equation for $\Omega(x, p)$, and then equation (6) becomes

$$\nabla_\xi \Omega = \left(2\xi_1 + 2\frac{p_2 \xi_1 \sqrt{5}}{(\xi_1^2 + \xi_2^2)^{1/2}} + 4p_1 - 2p_0,\ 2\xi_2 + 2\frac{p_2 \xi_1 \sqrt{5}}{(\xi_1^2 + \xi_2^2)^{1/2}} - 2p_1 - 2p_0,\ 2p_0 + 2 - 2\xi_3 \right).$$

Then $\nabla_\xi \Omega = 0$ can be solved to give

$$\xi = \left(p_0 - 2p_1 + p_2\sqrt{5}\left(1 + f(p)^2\right)^{-1/2},\ p_0 + p_1 - p_2\sqrt{5}\left(1 + f(p)^{-2}\right)^{-1/2},\ p_0 + 1 \right)$$

whenever $p_0 < 2p_1$, where $f(p) = (p_0 + p_1)/(p_0 - 2p_1)$. The resulting surface is show in Figure 3. The figure shows how $\Pi(V)$ forms a non-convex 2-dimensional surface which is contained within the convex hull.

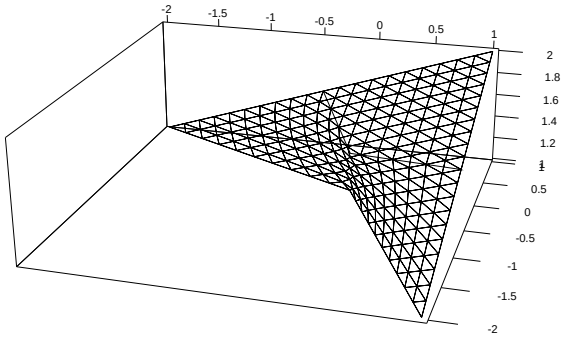

FIGURE 3. Persepective view of $\Pi(V)$ for the example in Section 3.4. The locus of the Fréchet mean is a 2-dimensional surface which resembles a rubber sheet pulled taut between the corners.

4. Projection onto the locus of the Fréchet mean and principal component analysis: algorithms

4.1. **Projection.** In order to use the surface $\Pi(V)$ as a principal component, we need to be able to project data onto $\Pi(V)$. Let $z \in \mathcal{T}_N$ denote a data point and fix $V = \{v_0, \ldots, v_k\}$. A projection of z onto $\Pi(V)$ is a point which minimizes $d(z, \Pi(V))$. This point might not be unique as $\Pi(V)$ is not convex. A naive algorithm to find a projection is to perform exhaustive search, as follows.

Algorithm 2. *Exhaustive search to project z onto $\Pi(V)$.*

1. *Construct a lattice of points $L \subset \mathcal{S}^k$. For $k = 2$ this is a triangular lattice.*
2. *For each point $p \in L$ use a standard algorithm to compute $\mu(V, p)$. (See Section 2.2.)*
3. *Find $p \in L$ which minimizes $d(z, \mu(V, p))$.*

We implemented this algorithm for $k = 2$ and used the algorithm of [3] at step 2 to compute Fréchet means. Algorithm 2 is computationally very expensive, since the resolution of the lattice L needs to be quite high in order to obtain accurate results. Consequently we only use the exhaustive search algorithm in what follows as a benchmark in order to assess other methods.

As an alternative to exhuastive search, we would like a more efficient algorithm which is defined entirely in terms of the geodesic geometry, since any reliance on local differentiable structure is likely to be problematic at orthant boundaries. We propose Algorithm 3, which we call the *geometric projection algorithm.*

Algorithm 3. *Geometric projection algorithm to project z onto $\Pi(V)$.*

Fix an initial estimate μ_0 of the projection of z, let $p = (0, \ldots, 0)$ and set $i = 0$.
***Repeat**:*

1. *Construct $\Gamma(\mu_i, v_j)$ for $j = 0, \ldots, k$.*
2. *For $j = 0, \ldots, k$ let $y_{i,j}$ be the point a proportion $s_i = 1/(i+2)$ along $\Gamma(\mu_i, v_j)$.*
3. *Find $r \in \{0, \ldots, k\}$ which minimizes $d(z, y_{i,r})$.*
4. *Set $\mu_{i+1} = y_{i,r}$ and let $p \leftarrow ip/(i+1) + e_r/(i+1)$ where e_r is the r-th*

standard basis vector in $\mathcal{S}^k$.
5. *Set* $i \leftarrow i + 1$.
Until *the sequence* $\mu_0, \mu_1, \ldots$ *converges.*

The algorithm is a modification of Sturm's algorithm for computing the Fréchet mean of V (Algorithm 1.) At each step of Sturm's algorithm, one of the points $y_{i,j}$ is used as the new estimate μ_{i+1}, and the point $y_{i,j}$ is sampled according to a fixed probability vector p. Here, the new estimate for the projection, μ_{i+1}, is again chosen from $y_{i,0}, \ldots, y_{i,k}$ but instead is selected to greedily minimize the distance from z. The vector $p \in \mathcal{S}^k$ estimates the weight vector associated to the projected point: at iteration i, $i \times p$ is a vector with integer entries which counts the number of times the algorithm has moved the estimate of the projection towards each vertex in V. The computational cost of the algorithm is similar to that for computing a single Fréchet mean using the Sturm algorithm. For $k = 2$ the initial point μ_0 is sampled uniformly from the perimeter of $\Pi(V)$. Convergence is tested as follows: at iteration i it is determined whether $d(\mu_s, \mu_t) < \epsilon$ for all $s, t \in \{i - m, \ldots, i\}$ where $\epsilon > 0$ and m are fixed. If that is the case then the algorithm terminates. The output from the algorithm after I iterations is an estimate μ_I of the projection of z and a vector $p \in \mathcal{S}^k$.

The geometric projection algorithm is presented here without a proof of convergence and without further theoretical study of its properties. Instead we rely on a simulation study in the next section to assess the effectiveness of the algorithm.

4.2. **Simulations.** We ran a set of simulations designed to demonstrate that, specifically in the case that $k = 2$, Algorithm 3 converges to a tree on $\Pi(V)$ which minimizes $d(z, \Pi(V))$. For each iteration of the simulation, a random species tree u with $N = 6$ taxa was generated under the [18] coalescent. Three trees v_0, v_1, v_2 and a fourth test tree z were then generated under a coalescent model constrained to be contained within the tree u, and thus corresponded to gene trees coming from the underlying species tree u. (See [25] for more information about the relationship between species trees and gene trees.) The DendroPy library [35] was used to generate these trees. The test tree z was then projected onto $\Pi(V)$ for $V = \{v_0, v_1, v_2\}$ using the exhaustive search algorithm and the geometric projection algorithm. All calculations were carried out ignoring pendant edges. This particular simulation scheme was chosen in order to generate a variety of different geometrical configurations for the points v_0, v_1, v_2, z, as well as being biologically reasonable. If the trees were sampled with topologies chosen independenly uniformly at random, for example, the simulation procedure would only have explored instances of $\Pi(V)$ with widely differing vertices.

The results obtained from the two algorithms were compared in two ways. First, the distances from the data tree to the trees obtained with the two algorithms were computed and checked to ensure that the projection algorithm obtained a distance less than or equal to the exhaustive search. Second, the distance between the tree from geometric projection and tree from exhaustive search was checked to ensure that the two trees were close together. For the second check we considered any distance greater than 1% of the total internal length of the data tree to be a failure.

In a run of 10,000 iterations of this procedure, 95.65% of the iterations passed the two tests. However, even the set of failing iterations produced a projection result which were quite close to the exhaustive search result. Among the 435 failing

iterations, the perpendicular distance for the projection was an average of 3.7% greater than the perpendicular distance of the exhaustive search, and the distance between the two results was an average of 4.7% of the total internal length of the data tree.

We believe that the failing results are attributable to the projection algorithm becoming trapped in local minima of the perpendicular distance. Starting the algorithm from several locations and comparing the results would help to mitigate this problem. However, for the present purpose of fitting higher-order principal components to a collection of data trees, we believe these small deviations from the exhaustive search solution are an acceptable trade for the great increase in computational speed obtained.

4.3. **Stochastic optimization for principal component analysis.** Given data $Z = \{z_1, \ldots, z_n\}$, our objective is to find $V = \{v_0, \ldots, v_k\}$ which minimizes the sum of squared projected distances $D_Z^2(\Pi(V))$. From this point on in the paper, we restrict to the case $k = 2$. The geometric projection algorithm is used to compute $D_Z^2(\Pi(V))$ given V, at least approximately, and so we must now consider how to search over the possible configurations of the vertices V. We adopt a stochastic optimization approach, Algorithm 4 below, which is similar to that for fitting principal geodesics in [27]. We assume we have available a set of proposals $M_1, \ldots, M_m$, each of which is a map from $\mathcal{T}_N$ to the set of distributions on $\mathcal{T}_N$. In particular, given any tree x, each $M_i(x)$ is asuumed to be a distribution on $\mathcal{T}_N$ from which we can easily sample.

Algorithm 4. *Stochastic optimization algorithm to fit* $\Pi(V)$ *to* Z*.*

Fix an initial set $V = \{v_0, v_1, v_2\}$ *and compute* $D_Z^2(\Pi(V))$*.*
Repeat*:*
  For $i = 0, 1, 2$*:*
   For $j = 1, \ldots, m$*:*
    1. *Sample a tree* w *from* $M_j(v_i)$*.*
    2. *Let* V' *be the set* V *but with* w *replacing* v_i*.*
    3. *Compute* $D_Z^2(\Pi(V'))$ *using the geometric projection algorithm.*
    4. *If* $D_Z^2(\Pi(V')) < D_Z^2(\Pi(V))$ *set* $V \leftarrow V'$*.*
Until *convergence.*

The optimization algorithm attempts to minimize $D_Z^2(\Pi(V))$ by stochastically varying one point $v \in V$ at a time using the proposals $M_i(v)$. The algorithm is greedy: whenever a configuration V' improves upon the current configuration V we replace V with V'. Convergence is assessed by considering the relative change in $D_Z^2(\Pi(V))$ over a certain fixed number of iterations. If this is less than some proportion then the algorithm terminates. We used three different types of proposal. The first samples a tree uniformly at random with replacement from the data set Z. The second type is a refinement of this: given a tree x it similarly samples a tree z uniformly at random with replacement from the data set Z. Then the geodesic $\Gamma(x, z)$ is computed, and a beta distribution is used to sample a tree some proportion of the distance along $\Gamma(x, z)$. The third type of proposal is a random walk starting from x, as described in [27]. The random walk proposals can have different numbers of steps and step-sizes. The algorithm is not guaranteed to find a global optimum, and it can become stuck in local minima. It is therefore necesary

to run the algorithm several times with different starting points for each data set, and then compare the results from each run.

Two statistics can be used to summarize the fit of $\Pi(V)$ to a data set Z: the sum of squared projected distances $D_Z^2(\Pi(V))$ and a non-Euclidean proportion of variance statistic, denoted r^2. If the projection of each data point z onto $\Pi(V)$ is denoted $\pi(z_i)$ and $\bar{\pi}$ denotes the Fréchet mean of $\pi(z_1), \ldots, \pi(z_n)$, then

$$r^2 = \frac{\sum_i d(z_i, \pi(z_i))^2}{\sum_i d(z_i, \pi(z_i))^2 + \sum_i d(\bar{\pi}, \pi(z_i))^2}.$$

The denominator in this expression varies with $\Pi(V)$ since Pythagoras' theorem does not hold in tree-space. Unlike $D_Z^2(\Pi(V))$, the r^2 statistic is quite sensitive to small changes in V, but it can be interpreted broadly as the proportion of variance explained by $\Pi(V)$.

In order to assess the performance of the algorithm we performed a small simulation study. Eight data sets of 100 trees containing $N = 10$ taxa were generated in the following way. For each data set a tree topology was sampled from a coalescent process, and each edge length was sampled from a gamma distribution with shape $\alpha = 2$ and rate $\beta = 20$, to give a tree w_0. Two trees w_1, w_2 were then obtained by applying random topological operations to w_0. In four of the data sets w_1, w_2 were obtained by performing nearest neighbour interchange operations, while in the other four data sets sub-tree prune and regraft operations were used. Then, to construct each data set given $W = \{w_0, w_1, w_2\}$, 100 points were sampled from a Dirichlet distribution on $\mathcal{S}^2$ with parameter $(4, 4, 4)$ and the corresponding points on $\Pi(W)$ were found using the Bačák algorithm. Each point was then perturbed by using a random walk, so that each data set resembled a cloud of points around the surface $\Pi(W)$. The step-size of the random walk was tuned to produce data sets classified as having either low or high dispersion. Table 1 summarizes the data sets used and the simulation results. It shows the sum of squared projected distances $D_Z^2(\Pi(W))$ (the 'true' sum of squared distances) and the fitted value $D_Z^2(\Pi(V))$, as well as the non-Euclidean r^2 statistic. The exhaustive projection algorithm was used to compute $D_Z^2(\Pi(W))$ while the geometric projection algorithm was used for $D_Z^2(\Pi(V))$. From the table it can be seen that the algorithm performs well in every scenario.

5. Results

5.1. **Coelacanths genome and transcriptome data.** We applied our proposed method to the dataset comprising 1,290 nuclear genes encoding 690,838 amino acid residues obtained from genome and transcriptome data by [22]. Over the last decades researchers have worked on the phylogenetic relations between coelacanths, lungfishes and tetrapods, but controversy remains despite several studies [13]. Most morphological and paleontological studies support the hypothesis that lungfishes are closer to tetrapods than they are to coelacanths (Tree 1 in Figure 1 from [22]). However, there exists research in the field that supports the hypothesis that coelacanths are closer to tetrapods (Tree 2 in Figure 1 from [22]). Others support the hypothesis that coelacanths and lungfishes form a sister clade (Tree 3 in Figure 1 from [22]) or tetrapods, lungfishes, and coelacanths cannot be resolved (Tree 4 in Figure 1 from [22]).

Topological scenario	Low dispersion D_Z^2	Low dispersion r^2	High dispersion D_Z^2	High dispersion r^2
2 NNI	0.284(0.270)	40.9%(49.8%)	2.65(2.73)	17.5%(18.0%)
4 NNI	0.310(0.298)	61.4%(65.9%)	2.57(2.91)	26.5%(20.4%)
2 SPR	0.255(0.254)	58.6%(61.8%)	2.17(2.41)	28.5%(20.7%)
4 SPR	0.269(0.278)	54.0%(48.2%)	2.39(2.78)	24.3%(21.9%)

TABLE 1. Simulations to assess the stochastic optimization algorithm. The left column describes number and type of topological operation used to obtain w_1, w_2 from w_0 for each data set. The abbreviation NNI stands for nearest neighbour interchange and SPR stands for sub-tree prune and regraft. For each scenario, two data sets were generated by perturbing points on $\Pi(W)$ via random walks, with low and high dispersions respectively. The table shows the value of $D_Z^2(\Pi(V))$ for the fitted principal component and in brackets the reference value $D_Z^2(\Pi(W))$. Similarly the non-Euclidean r^2 statistic is shown with the reference value in brackets.

We reconstructed gene trees using the R package "Phangorn" [32], and each gene tree was estimated using the maximum likelihood (optim.pml and pml functions) under the Le-Gascuel (LG) model [20]. The data set consisted of 1290 gene alignments for 10 species. The species were lungfish (*Protopterus annectens*, denoted Pa), coelacanth (*Latimeria chalumnae*, Lc), and three tetrapods: frog (*Xenopus tropicalis*, Xt), chicken (*Gallus gallus*, Gg) and human (*Homo sapiens*, Hs). Two ray-finned fish, *Danio rerio* (denoted Dr) and *Takifugu rubripes* (denoted Tr), along with three cartilaginous fish (*Scyliorhinus canicula, Leucoraja erinacea, Callorhinchus milii*) were included as an out-group. The cartilaginous fish will be denoted Sc, Le, and Cm respectively.

Analysis was performed ignoring pendant edge lengths. A total of 97 outlying trees were removed using KDETrees [36], so that 1193 gene trees remained. The Fréchet mean was computed using the Bačák algorithm and its topology is shown in Figure 4. The mean tree does not resolve whether coelacanth or lungfish is the closest relative of the tetrapods. The sum of squared distances of the data points to the Fréchet mean was 19.7. A principal geodesic was constructed using the algorithm from [27]: the sum of squared projected distances was 9.53 and the r^2 statistic was 51.4%. Traversing the principal geodesic gives trees with the same topology as the Fréchet mean which contract down to a star tree at one end of the geodesic, and expand in size at the other end. This shows that the principal source of variation in the data set is the overall scale of the gene trees, or in other words, the total amount of evolutionary divergence for each gene.

Figure 4 illustrates the second principal component. The sum of squared projected distances was 7.29 and the r^2 statistic was 61.8%. This represents a relatively small increase in the proportion of variance in relation to the principal geodesic. Three runs of Algorithm 4 were performed to construct the second principal component. The results obtained had very similar summary statistics, but the topologies

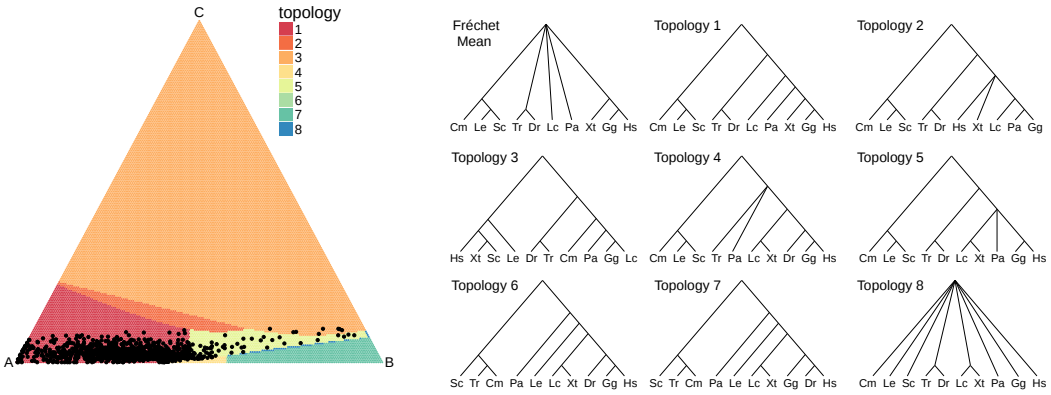

FIGURE 4. The second principal component computed from the lungfish data set. Left: The simplex shaded according to the topology of the corresponding points on $\Pi(V)$. The projections of the data points are also displayed. Right: topologies of trees on $\Pi(V)$.

displayed on the surfaces were more variable. Figure 4 is therefore a representative choice. Although the projected points are clustered towards the bottom of the simplex in the figure, the full simplex was drawn to show all the different topological regions. The points can be separated by zooming in on the simplex. Of the 1193 gene trees, 1094 projected to points with topology 1, which supports lungfish as the closest relative of the tetrapods (Pa grouped next to Hs, Gg, Xt). From the remaining projected data points, 75 have topology 5. This topology places both lungfish and coelacanth in a clade with the tetrapods. Several topologies (3,4,6 and 7) have biologically implausible relationships. However, the projected data points lying outside topology 1 all lie close to the boundary of their respective orthants (with at least one edge length less than 0.0005), so for example, the projected data points with topology 3 have very short edge lengths for the biologically implausible clades (such as the grouping of Xt with Sc) and so lie close to trees with more plausible topologies. Overall, the second principal component suggests that the data support lungfish as the closest relative of tetrapod (topology 1), and that most of the variation within the data comes from edge length variation within that topology rather than from conflicting topologies. It is interesting to note that the Fréchet mean and principal geodesic did not exhibit topology 1, and that the second order principal component was needed to resolve the controversial relationship between the coelacanth, lungfish and tetrapods. The exhaustive projection algorithm was used to project the data onto the surface $\Pi(V)$ produced by Algorithm 4, in order to compare with the results obtained by geometric projection. The sum of squared distances between the projected trees obtained with the two different algorithms was 0.004, a small fraction of the sum of squared projected distances 7.29 for $\Pi(V)$.

5.2. **Apicomplexa.** We also applied our method to a set trees constructed from of 268 orthologous sequences from eight species of protozoa presented in [19]. The data set from [19] consists of gene trees reconstructed from the following sequences: *Babesia bovis* (Bb) [5] from GenBank (GenBank accession numbers

AAXT01000001–AAXT01000013), *Cryptosporidium parvum* (Cp) [1] from CryptoDB.org [14], *Eimeria tenella* (Et) from GeneDB.org [15], *Plasmodium falciparum* (Pf) [11] and *Plasmodium vivax* (Pv) from PlasmoDB.org [2], *Theileria annulata* (Ta) [30] from GeneDB.org [15], and *Toxoplasma gondii* (Tg) from Toxo-DB.org [10]. A free-living ciliate, *Tetrahymena thermophila* (Tt) [8], was used as the outgroup.

The phylum Apicomplexa contains many important protozoan pathogens [21], including the mosquito-transmitted *Plasmodium* spp., the causative agents of malaria; *T. gondii*, which is one of the most prevalent zoonotic pathogens worldwide; and the water-born pathogen *Cryptosporidium* spp. Several members of the Apicomplexa also cause significant morbidity and mortality in both wildlife and domestic animals. These include *Theileria* spp. and *Babesia* spp., which are tick-borne haemoprotozoan ungulate pathogens, and several species of *Eimeria*, which are enteric parasites that are particularly detrimental to the poultry industry. Due to their medical and veterinary importance, whole genome sequencing projects have been completed for multiple prominent members of the Apicomplexa. We removed 16 outlier trees using the KDETrees software [36] before fitting principal components.

The trees were analysed ignoring pendant edges. The Fréchet mean was computed using the Bačák algorithm: the corresponding tree topology was unresolved, and is shown in Figure 5. The sum of squared distances from the mean to the data points was 24.6. The principal geodesic was estimated using the algorithm from [27]. The principal geodesic has a non-euclidean r^2 score of 40% and the sum of squared projected distances was 14.2. The principal geodesic displays two main effects: (i) the edges leading to the (Pv, Pf) clade, (Tg, Et) clade and (Bb, Ta) clade vary substantially in length and (ii) a topological re-arrangement whereby the clade containing (Pv, Pf) paired with (Et, Tg) is replaced with a clade containing (Pv, Pf) paired with (Bb, Ta). However, the second effect involved very short internal edges, so that along its length, the trees on the principal geodesic resembled the mean tree shown in Figure 5 but with different overall scale. The principal geodesic therefore reflects variation in the scale of the tree.

Figure 5 illustrates the second principal component, with the simplex shaded according to the corresponding tree topology on $\Pi(V)$. Three separate runs of Algorithm 4 converged to give similar results. The summary statistics for the second principal component are: sum of squared projected distances 10.3, r^2 statistic 56%. While these summary statistics were consistent between runs, the set of topologies displayed on $\Pi(V)$ was subject to more variation, so Figure 5 is a representative choice, although topologies 1,4 and 6 were present in all runs. The results show how the second principal component is able to tease more from the data than the variation in overall scale captured by the principal geodesic. Topology 4 is congruent with the generally accepted phylogeny of taxa within the Apicomplexa and is a resolution of the Fréchet mean tree: *Theileria annulata* (Ta) and *Babesia bovis* (Bb) group together; the two Plasmodium species (Pf and Pv) group together; *Cryptosporidium parvum* is the deepest rooting apicomplexan; and Pv, Pf, Ta, Bb are monophyletic (they are all hemosporidians or blood parasites). The figure shows that the second principal component corresponds to variation in topology consisting of nearest neighbour interchange operations which transform topology 4 into topologies 1 and 6. None of the projected trees have topology 5, although this is the topology of one of the vertices of $\Pi(V)$. This topology appears to be present

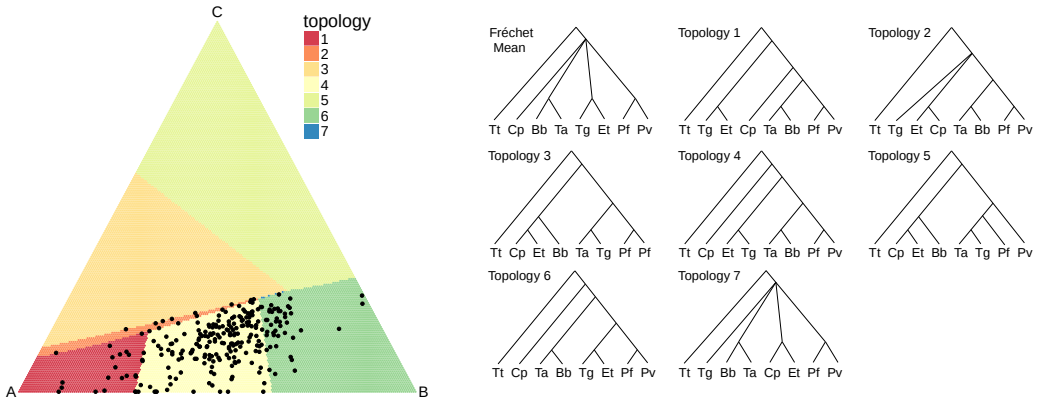

FIGURE 5. The second principal component computed from apicomplexa data set. Left: The simplex shaded according to the topology of the corresponding points on $\Pi(V)$. The projections of the data points are also displayed. Right: topologies of trees on $\Pi(V)$.

in order for $\Pi(V)$ to be positioned in such a way as to capture the other topologies. Topology 2 shows evidence of stickiness (see Section 3.1): although the topology is unresolved, so that the coloured triangle lies in a codimension-1 region of tree-space, it occupies non-zero area on the simplex. As for the lungfish, the exhaustive and geometric projection algorithms were compared on the surface $\Pi(V)$ produced by Algorithm 4. The distances between the projected points obtained with the two algorithms were very small compared to the distances of the data points from $\Pi(V)$ (the sum of squared distances between pairs of projected points was 3.91E-4).

6. Discussion

This paper presents three main innovations: (i) use of the locus of the Fréchet mean $\Pi(V)$ as an analogue of a principal component in tree-space, (ii) proof that $\Pi(V)$ has the desired dimension, and (iii) the geometric projection algorithm for projecting data onto $\Pi(V)$. The locus of the Fréchet mean was first proposed as a geometric object for principal component analysis in tree-space by [38], though in [31] Pennec has made a similar proposal for analogues of principal component analysis in Riemannian manifolds and other geodesic metric spaces. The barycentric subspaces of Pennec correspond exactly to the surfaces $\Pi(V)$ considered in this paper. Pennec's methodology, however, is principally based in the context of a Riemannian manifold rather than in tree-space, though he points out the potential for generalization. There are substantial differences between Pennec's analogue of principal component analysis on Riemannian manifolds (barycentric subspace analysis) and the methodology presented in this paper. In particular, a key aim of barycentric subspace analysis is to produce *nested* principal components, in the sense of equation (2), while we do not make that restriction here. For example, if we consider a surface $\Pi(V)$ for $k = 2$ then the only geodesics which are obviously

contained in $\Pi(V)$ are the edges, and it is unappealing to restrict the principal geodesic to be one of these.

Our analysis has been restricted to data sets with relatively few taxa and to the construction of the first and second principal components. The algorithms presented in this paper scale linearly with respect to the number of data points n, but run in polynomial time with respect to the number of taxa N. However, by partitioning the data set for the geometric projection algorithm, parallel computer architectures can be employed and the speed-up is approximately proportional to the number of processors used. While the geometric projection algorithm runs relatively quickly, the calculations involved in searching for the optimal set of vertices V can be very substantial. The experimental data sets in Section 5 took between 1 and 3 days to analyse, running on 4 processors each. For higher order components with $k > 2$, this computational burden will increase, and it is likely that finding a global minimum for $D_Z^2(\Pi(V))$ will be more difficult. The figures in Section 5 show the potential for creating visualizations of the data which reveal meaningful biological structure. The pattern of projected points obtained for the experimental data sets we considered were very similar to results obtained via multidimensional scaling. However, multidimensional scaling is not capable of revealing the features of the data set which cause the observed variation. More information could be included in the graphical representation of our results, such as the distance of the data points from their projections, information about the principal geodesic, and the proximity of points to orthant boundaries.

We presented the geometric projection algorithm without a proof of convergence, and used simulation to assess its accuracy. The algorithm is attractive in that it is defined entirely in terms of the geodesic structure on tree-space, and so it could be used on any geodesic metric space, including Riemannian manifolds. The algorithm clearly deserves further investigation, and we aim to study its properties in a future publication.

Acknowledgement

The authors thank D. Howe from University of Kentucky for useful comments on the analysis of apicomplexa data set.

References

[1] Mitchell S. Abrahamsen, Thomas J. Templeton, Shinichiro Enomoto, Juan E. Abrahante, Guan Zhu, Cheryl A. Lancto, Mingqi Deng, Chang Liu, Giovanni Widmer, Saul Tzipori, Gregory A. Buck, Ping Xu, Alan T. Bankier, Paul H. Dear, Bernard A. Konfortov, Helen F. Spriggs, Lakshminarayan Iyer, Vivek Anantharaman, L. Aravind, and Vivek Kapur. Complete genome sequence of the apicomplexan, cryptosporidium parvum. *Science*, 304:441–445, 2004.

[2] Amit Bahl, Brian Brunk, Jonathan Crabtree, Martin J. Fraunholz, Bindu Gajria, Gregory R. Grant, Hagai Ginsburg, Dinesh Gupta, Jessica C. Kissinger, Philip Labo, Li Li, Matthew D. Mailman, Arthur J. Milgram, David S. Pearson, David S. Roos, Jonathan Schug, Christian J. Stoeckert, and Patricia Whetzel. Plasmodb: the plasmodium genome resource. a database integrating experimental and computational data. *UCLA Sym. Bi.*, 31:212–215, 2003.

[3] M. Bačák. Computing medians and means in Hadamard spaces. *SIAM J. Optimiz.*, 24(3):1542–1566, 2014.

[4] L. Billera, S. Holmes, and K. Vogtman. Geometry of the space of phylogenetic trees. *Adv. Appl. Math*, 27:733–767, 2001.

[5] Kelly A. Brayton, Audrey O. T. Lau, David R. Herndon, Linda Hannick, Lowell S. Kappmeyer, Shawn J. Berens, Shelby L. Bidwell, Wendy C. Brown, Jonathan Crabtree, Doug Fadrosh, Tamara Feldblum, Heather A. Forberger, Brian J. Haas, Jeanne M. Howell, Hoda

Khouri, Hean Koo, David J. Mann, Junzo Norimine, Ian T. Paulsen, Diana Radune, Qinghu Ren, Roger K. Smith Jr., Carlos E. Suarez, Owen White, Jennifer R. Wortman, Donald P. Knowles Jr.1, Terry F. McElwain, and Vishvanath M. Nene. Genome sequence of babesia bovis and comparative analysis of apicomplexan hemoprotozoa. *PLOS Pathog.*, 3:e148, 2007.

[6] M. R. Bridson and A. Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319. Springer-Verlag, 2011.

[7] Chris Ding and Xiaofeng He. K-means clustering via principal component analysis. In *Proceedings of the twenty-first international conference on Machine learning*, page 29. ACM, 2004.

[8] Jonathan A Eisen, Robert S Coyne, Martin Wu, Dongying Wu, Mathangi Thiagarajan, Jennifer R Wortman, Jonathan H Badger, Qinghu Ren, Paolo Amedeo, Kristie M Jones, Luke J Tallon, Arthur L Delcher, Steven L Salzberg, Joana C Silva, Brian J Haas, William H Majoros, Maryam Farzad, Jane M Carlton, Roger K Smith, Jr., Jyoti Garg, Ronald E Pearlman, Kathleen M Karrer, Lei Sun, Gerard Manning, Nels C Elde, Aaron P Turkewitz, David J Asai, David E Wilkes, Yufeng Wang, Hong Cai, Kathleen Collins, B. Andrew Stewart, Suzanne R Lee, Katarzyna Wilamowska, Zasha Weinberg, Walter L Ruzzo, Dorota Wloga, Jacek Gaertig, Joseph Frankel, Che-Chia Tsao, Martin A Gorovsky, Patrick J Keeling, Ross F Waller, Nicola J Patron, J. Michael Cherry, Nicholas A Stover, Cynthia J Krieger, Christina del Toro, Hilary F Ryder, Sondra C Williamson, Rebecca A Barbeau, Eileen P Hamilton, and Eduardo Orias. Macronuclear genome sequence of the ciliate *Tetrahymena thermophila*, a model eukaryote. *PLOS Biol.*, 4:1620–1642, 2006.

[9] A. Feragen, M. Owen, J. Petersen, M.M.W. Wille, L.H. Thomsen, A. Dirksen, and de Bruijne M. Tree-space statistics and approximations for large-scale analysis of anatomical trees. In *23rd biennial International Conference on Information Processing in Medical Imaging (IPMI)*, 2013.

[10] Bindu Gajria, Amit Bahl, John Brestelli, Jennifer Dommer, Steve Fischer, Xin Gao, Mark Heiges, John Iodice, Jessica C. Kissinger, Aaron J. Mackey, Deborah F. Pinney, David S. Roos, Christian J. Stoeckert, Haiming Wang, and Brian P. Brunk. Toxodb: an integrated toxoplasma gondii database resource. *UCLA Sym. Bi.*, 36:D553–D556, 2008.

[11] Malcolm J. Gardner, Neil Hall, Eula Fung, Owen White, Matthew Berriman, Richard W. Hyman, Jane M. Carlton, Arnab Pain, Karen E. Nelson, Sharen Bowman, Ian T. Paulsen, Keith James, Jonathan A. Eisen, Kim Rutherford, Steven L. Salzberg, Alister Craig, Sue Kyes, Man-Suen Chan, Vishvanath Nene, Shamira J. Shallom, Bernard Suh, Jeremy Peterson, Sam Angiuoli, Mihaela Pertea, Jonathan Allen, Jeremy Selengut, Daniel Haft, Michael W. Mather, Akhil B. Vaidya, David M. A. Martin, Alan H. Fairlamb, Martin J. Fraunholz, David S. Roos, Stuart A. Ralph, Geoffrey I. McFadden, Leda M. Cummings, G. Mani Subramanian, Chris Mungall, J. Craig Venter, Daniel J. Carucci, Stephen L. Hoffman, Chris Newbold, Ronald W. Davis, Claire M. Fraser, and Bart Barrell. Genome sequence of the human malaria parasite plasmodium falciparum. *Nature*, 419:498–511, 2002.

[12] Kevin Gori, Tomasz Suchan, Nadir Alvarez, Nick Goldman, and Christophe Dessimoz. Clustering genes of common evolutionary history. *Mol. Biol. Evol.*, page msw038, 2016.

[13] S.B. Hedges. Vertebrates (vertebrata). *In: Hedges SB, Kumar S, editors. The timetree of life*, pages 309–314, 2009.

[14] Mark Heiges, Haiming Wang, Edward Robinson, Cristina Aurrecoechea, Xin Gao, Nivedita Kaluskar, Philippa Rhodes, Sammy Wang, Cong-Zhou He, Yanqi Su, John Miller, Eileen Kraemer, and Jessica C. Kissinger. Cryptodb: a cryptosporidium bioinformatics resource update. *UCLA Sym. Bi*, 34:D419–D422, 2006.

[15] Christiane Hertz-Fowler, Chris S. Peacock, Valerie Wood, Martin Aslett, Arnaud Kerhornou, Paul Mooney, Adrian Tivey, Matthew Berriman, Neil Hall, Kim Rutherford, Julian Parkhill, Alasdair C. Ivens, Marie-Adele Rajandream, and Bart Barrell. Genedb: a resource for prokaryotic and eukaryotic organisms. *UCLA Sym. Bi.*, 32:D339–D343, 2004.

[16] D.M. Hillis, T.A. Heath, and K. St. John. Analysis and visualization of tree space. *Syst. Biol.*, 54(3):471–482, 2005.

[17] T. Hotz, S. Huckemann, H. Le, J.S. Marron, J.C. Mattingly, E. Miller, J. Nolen, M. Owen, V. Patrangenaru, and S. Skwerer. Sticky central limit theorems on open books. *Ann. Appl. Probab.*, 23(6):2238–2258, 2013.

[18] John Frank Charles Kingman. The coalescent. *STOCH PROC APPL*, 13(3):235–248, 1982.

[19] C. Kuo, J. P. Wares, and J. C. Kissinger. The apicomplexan whole-genome phylogeny: An analysis of incongruence among gene trees. *Mol. Biol. Evol.*, 25(12):2689–2698, 2008.
[20] S.Q. Le and O. Gascuel. An improved general amino acid replacement matrix. *Mol. Biol. Evol.*, 25(7):1307–1320, 2008.
[21] ND Levine. Progress in taxonomy of the apicomplexan protozoa. *J. Eukaryot. Microbiol.*, 35:518–520, 1988.
[22] D. Liang, X.X. Shen, and P. Zhang. One thousand two hundred ninety nuclear genes from a genome-wide survey support lungfishes as the sister group of tetrapods. *Mol. Biol. Evol.*, 30(8):1803–1807, 2013.
[23] B. Lin, B. Sturmfels, X. Tang, and R. Yoshida. Convexity in tree spaces. *arXiv:1510.08797*, 2015.
[24] Anna Lubiw, Daniela Maftuleac, and Megan Owen. Shortest paths and convex hulls in 2d complexes with non-positive curvature. *arXiv preprint arXiv:1603.00847*, 2016.
[25] Wayne P Maddison. Gene trees in species trees. *Syst. Biol.*, 46(3):523–536, 1997.
[26] E. Miller, M. Owen, and J. S. Provan. Polyhedral computational geometry for averaging metric phylogenetic trees. *Adv. Appl. Math.*, 68:51–91, 2015.
[27] T.M.W. Nye. An algorithm for constructing principal geodesics in phylogenetic treespace. *IEEE ACM T. Comput. Bi.*, 11(2):304–315, 2014.
[28] Tom M. Nye. Principal components analysis in the space of phylogenetic trees. *Ann. Stat.*, 39(5):2716–2739, 2011.
[29] M. Owen and J. S. Provan. A fast algorithm for computing geodesic distances in tree space. *IEEE ACM T. Comput. Bi.*, 8(1):2–13, 2011.
[30] A. Pain, H. Renauld, and et al. Genome of the host-cell transforming parasite theileria annulata compared with t. parva. *Science*, 309:131–133, 2005.
[31] Xavier Pennec. Barycentric subspace analysis on manifolds. *arXiv preprint arXiv:1607.02833*, 2016.
[32] K.P. Schliep. Phangorn: Phylogenetic analysis in R. *Bioinformatics*, 27(4):592–593, 2011.
[33] C. Semple and M. Steel. *Phylogenetics, Oxford Lecture Series in Mathematics and its Applications, 24*. Oxford University Press, 2003.
[34] Karl-Theodor Sturm. Probability measures on metric spaces of nonpositive curvature. *Heat Kernels and Analysis on Manifolds, Graphs, and Metric Spaces: Lecture Notes from a Quarter Program on Heat Kernels, Random Walks, and Analysis on Manifolds and Graphs: April 16-July 13, 2002, Emile Borel Centre of the Henri Poincaré Institute, Paris, France*, 338:357, 2003.
[35] J. Sukumaran and Mark T. Holder. Dendropy: A python library for phylogenetic computing. *Bioinformatics*, 26:1569–1571, 2010.
[36] G. Weyenberg, R. Yoshida, and D. Howe. Normalizing kernels in the Billera-Holmes-Vogtmann treespace. *IEEE ACM T. Comput. Bi.*, page doi:10.1109/TCBB.2016.2565475, 2016.
[37] Grady Weyenberg, Peter M Huggins, Christopher L Schardl, Daniel K Howe, and Ruriko Yoshida. KDEtrees: non-parametric estimation of phylogenetic tree distributions. *Bioinformatics*, 30(16):2280–2287, 2014.
[38] Grady S Weyenberg. *Statistics in the Billera-Holmes-Vogtmann Treespace*. PhD thesis, University of Kentucky, 2015.
[39] H. Zha, C. Ding, M. Gu, X. He, and H.D. Simon. Spectral relaxation for k-means clustering. *Neu. Inf. Pro.*, 14:1057–1064, 2001.

School of Mathematics and Statistics, Newcastle University, NE1 7RU, U.K.
E-mail address: tom.nye@ncl.ac.uk

Faculty 3: Mathematics / Computer Sciences, University of Bremen, Bremen 28359, Germany
E-mail address: xtang@uni-bremen.de

MRC Integrative Epidemiology Unit, University of Bristol, Oakfield House, Oakfield Grove, Bristol, BS8 2BN, U.K.
E-mail address: grady.weyenberg@bristol.ac.uk

Department of Operations Research, Naval Postgraduate School, Monterey, CA 93943-5219, U.S.A.
E-mail address: ryoshida@nps.edu